

Deep Learning For Video Saliency Detection

Wenguan Wang, and Jianbing Shen, *Senior Member, IEEE*, and Ling Shao, *Senior Member, IEEE*

Abstract—This paper proposes a deep learning model to efficiently detect salient regions in videos. It addresses two important issues: (1) deep video saliency model training with the absence of sufficiently large and pixel-wise annotated video data; and (2) fast video saliency training and detection. The proposed deep video saliency network consists of two modules, for capturing the spatial and temporal saliency stimuli, respectively. The dynamic saliency model, explicitly incorporating saliency estimates from the static saliency model, directly produces spatiotemporal saliency inference without time-consuming optical flow computation. We further propose a novel data augmentation technique that simulates video training data from existing annotated image datasets, which enables our network to learn diverse saliency stimuli and prevents overfitting with the limited number of training videos. Leveraging our synthetic video data (150K video sequences) and real videos, our deep video saliency model successfully learns both spatial and temporal saliency stimuli, thus producing accurate spatiotemporal saliency estimate. We advance the state-of-the-art on the DAVIS dataset (MAE of .06) and the FBMS dataset (MAE of .07), and do so with much improved speed (2fps with all steps) on one GPU.

Index Terms—Video saliency, deep learning, synthetic video data, salient object detection, convolutional neural network.

I. INTRODUCTION

Saliency detection has recently attracted a great amount of research interest. The reason behind this growing popularity lies in the effective use of these models in various vision tasks, such as image segmentation, object detection, video summarization and compression, to name a few. Saliency models can be broadly classified into two categories: human eye fixation prediction or salient object detection. According to the type of input, they can be further categorized into static and dynamic saliency models. While static models take still images as input, dynamic models work on video sequences. In this paper, we focus on detecting distinctive regions in dynamic scenes. *Convolutional neural networks* (CNNs) have been successfully utilized in many fundamental areas of computer vision, including object detection [1], [2], semantic segmentation [3], and still saliency detection [4], [5]. Inspired by this, we investigate CNNs to another computer vision task, namely video saliency detection.

The first problem of applying CNNs to video saliency is the lack of sufficiently large, densely labelled video training data. As far as we know, the successes of CNNs in computer vision are largely attributed to the availability of large-scale annotated images (e.g., ImageNet [6]). However, existing video datasets are too small to provide adequate training data for CNNs.

W. Wang and J. Shen are with Beijing Laboratory of Intelligent Information Technology, School of Computer Science, Beijing Institute of Technology, Beijing 100081, P. R. China. (email: shenjianbing@bit.edu.cn)

L. Shao is with the School of Computing Sciences, University of East Anglia, Norwich NR4 7TJ, U.K. (email: ling.shao@ieee.org)

TABLE I
STATISTICS FOR IMAGENET [6], FBMS [7], SEGTRACKV2 [8],
VSB100 [9] AND DAVIS [10] DATASETS.

Dataset	Clips	Annotations (\approx)
ImageNet	-	1.3×10^6
FBMS	59	500
SegTrackV2	14	1500
VSB100	100	600
DAVIS	50	4000

In Table 1, we list the statistics of the ImageNet dataset and widely adopted video object segmentation datasets, including FBMS [7], SegTrackV2 [8], VSB100 [9] and DAVIS [10]. It can be observed that, compared with popular image segmentation datasets like ImageNet, the existing video datasets are far from the requirement of training CNNs for video processing, in either quality or quantity. On the other hand, for the moment, creating such a large-scale video dataset is usually infeasible, because annotating videos is complex and time-consuming. To this end, we propose a video data augmentation approach to synthetically generating labeled video training data, which explicitly leverages existing large-scale image segmentation datasets. The simulated video data are easily accessible and rapidly generated, close to realistic video sequences and present various motion patterns, deformations, accompanied with automatically generated annotations and optical flow. The experimental results via these automatically generated videos clearly demonstrate the practicability of our strategy.

Our video data synthesis approach clears the underlying challenge for learning CNNs for many applications in video processing, where dynamic saliency detection is of no exception. Another challenge for detecting saliency in dynamic scenarios derives from the natural demand of this task. As suggested by human visual perception research [11], [12], when computing dynamic saliency maps, video saliency models need to consider both the spatial and the temporal characteristics of the scene. We propose a deep video saliency model for producing spatiotemporal saliency via fully exploring both the static and dynamic saliency stimuli. The proposed model adopts fully convolutional networks (FCNs) [3] for pixel-wise saliency prediction. For learning dynamic saliency stimulus, the proposed deep video saliency model learns from a large number of labelled videos, including both human-generated and natural video data, in a supervised learning mode. Associated with existing rich image saliency data, the static saliency stimuli is deeply guided and explicitly encoded in the above learning process via transferring and fine-tuning recent success in image classification [13].

Another important contribution of this work is that our deep video saliency model is much more computationally efficient

compared with existing video saliency models. Salient object detection is a key step in many image analysis tasks as it not only identifies relevant parts of a visual scene but may also reduce computational complexity by filtering out irrelevant segments of the scene. In recent years, some notable video saliency models have been proposed and show usefulness in many computer vision applications, such as video segmentation [14] and video re-timing [15]. However, time efficiency becomes the common major bottleneck for the applicability of existing video saliency algorithms; most computation time has been spent for optical flow computation. Additionally, from the perspective of learning deep networks in dynamic scenes, many schemes [16], [17], [18] take optical flow as input, causing high computational expenses.

In this work, we propose a both effective and efficient video saliency model, which frees itself from the computationally expensive optical flow estimation. One of the key insights of this paper is that, unlike high-level video applications such as action detection, video saliency can derive from short-term analysis of video frames. Thus we directly capture temporal saliency via learning deep networks from frame pairs, instead of using long-term video information, such as optical flows from multiple adjacent video frames.

We comprehensively evaluate our method on the FBMS dataset [7], where the proposed video saliency model produces more accurate saliency maps than state-of-the-arts. Meanwhile, it achieves a frame rate of 2fps (including all steps) on a GPU. Thus it is a practical video saliency detection model in terms of both speed and accuracy. We also report results on the newly released DAVIS dataset [10] and observe performance improvements over current competitors.

To summarize, the main contributions of this paper are threefold:

- We investigate convolutional neural networks for salient object detection in dynamic scenes. We adopt fully convolutional networks for end-to-end training and pixel-wise saliency prediction.
- We propose a novel training scheme based on synthetically generated video data, which explicitly leverages existing rich image datasets; both static and dynamic saliency stimuli are encoded into a unified deep learning model.
- Our methods are computationally efficient, much faster than traditional video saliency models and other deep networks in dynamic scenes.

The rest of this paper is structured as follows: An overview of the related work is given in Section II. Section III defines our proposed deep saliency model. The proposed synthetic video generation approach is articulated in Section IV. Section V shows experiment results on different databases and compare with the state-of-the-art methods. Finally, concluding remarks can be found in Section VI.

II. RELATED WORK

In this section, we give a brief overview of recent works in two lines: saliency detection, and deep learning models in dynamic scenes.

A. Saliency Detection

The saliency detection has been extensively studied in computer vision, and saliency detection models in general can be categorized into visual attention prediction or salient object detection. The former methods [11], [19], [20], [21] try to predict scene locations where a human observer may fixate. Salient object detection [22], [23], [24] aims at uniformly highlighting the salient regions, which has been shown benefit to a wide range of computer vision applications. More detailed reviews of the saliency models can be found in [25], [26].

Saliency models can be further divided into static and dynamic ones according to their input. In this work, we aim at detecting saliency object regions in videos.

Image saliency detection has been extensively studied for decades and a variety of image saliency methods have been proposed driven by the well-known bottom-up strategy. Numerous bottom-up models [22], [23] have been proposed based on detecting *contrast*, assuming salient regions in the visual field would first pop out from their surroundings and computing feature-based contrast followed by various mathematical principles. Meanwhile, some other mechanisms [24], [27], [28] have been proposed to adopt some prior knowledge, such as *background prior*, or global information, to detect salient objects in still images. More recently, deep learning techniques have been introduced to image saliency detection. These methods [4], [29] typically use CNNs to examine a large number of region proposals, from which the salient objects are selected. Currently, more and more methods [30] tend to learn in an end-to-end manner and directly generate pixel-wise saliency maps via fully convolutional networks (FCNs) [3], [31].

Compared with saliency detection in still images, detecting saliency in videos is a much more challenging problem due to the complication in the detection and utilization of temporal and motion information. So far, only a limited number of algorithms have been proposed for spatiotemporal saliency detection for video signals. The early models of dynamic saliency [32], [33], [34] can be viewed as simple extensions of exiting static saliency models with extra temporal dimension. Some more recent and notable approaches [14], [35], [36], [37], [38] to this task have been proposed, showing inspired performance and good potentials in many computer vision applications. However, the applicability of these approaches is severely limited by their high-computational costs. The main computational bottleneck comes from optical flow estimation, which contributes much to the promising results.

In this paper, we propose a deep learning model for detecting salient object regions in videos. The proposed model is effective, yet more computationally efficient compared with existing video saliency models. To the best of our knowledge, this paper is the first work that aims to detect video saliency using end-to-end FCNs.

B. Deep Learning Models in Dynamic Scenes

In this section, we mainly focus on famous, deep learning models for computer vision applications in dynamic scenes, including action recognition [16], [39], object segmentation

[40], [18], object tracking [41], [42], [43], [44], [45], attention prediction [17] and semantic segmentation [46], and explore their architectures and training schemes. This will help to clarify how our approach differs from previous efforts and will help to highlight the important benefits in terms of effectiveness and efficiency.

Many approaches [41], [42], [46] directly feed single video frames into neural networks trained on image data and adopt various techniques for post-processing the results with temporal or motion information. Unfortunately, these neural networks give up learning the temporal information which is often very important in video processing applications.

A famous architecture for training CNNs for action recognition in videos is proposed in [16], which incorporates two-stream convolutional networks for learning complementary information on appearance and motion. Other works [17], [40] adopt this architecture for dynamic attention prediction and video object segmentation. However, these methods train their models on multi-frame dense optical flow, which causes heavy computational burden.

In the areas of human pose estimation and video object processing, online learning strategy is introduced for improving performance [18], [39], [43], [44], [45]. Before processing an input video, these approaches generate various training samples for fine-tuning the neural networks learned from image data, thus enabling the models to be optimized towards the object of interest in the test video sequence. Obviously, these models are typically time-consuming as training neural networks is a slow process. Besides, the learned models are only specialized for specific classes of objects.

In this work, we show the possibilities of learning to detect generic salient objects in dynamic scenes by training on videos and images via an entirely offline manner. We proposed a novel technique for synthesizing video data via leveraging large amounts of image training data. The CNNs model can be efficiently and entirely trained on rich video sequences and images, thus successfully learning both static and dynamic saliency stimuli. Meanwhile, it directly learning inner relationship between frames, getting rid of time-consuming motion computation. Thus, our algorithm is significantly faster than traditional video saliency methods and the deep learning architectures that demand optical flow as input. In summary, our CNNs model learns to detect video saliency in a fast and effective manner.

III. DEEP NETWORKS FOR VIDEO SALIENCY DETECTION

In this work, we describe a procedure for constructing and learning deep video saliency networks using a novel synthetic video data generation approach. Our approach generates a large amount of video data (200K paired frames) from existing image datasets, and associates these annotated video sequences with existing labeled video and image data to learn deep video saliency networks. We first introduce the proposed CNNs based video saliency model in this section and then we describe our video synthesis approach in Sec. IV.

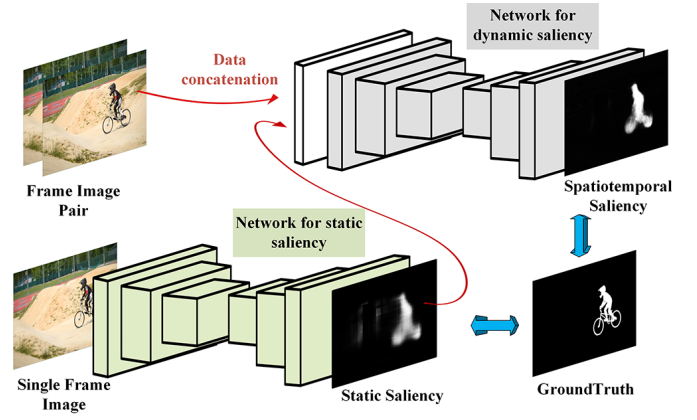


Fig. 1. A schematic representation of our proposed deep video saliency model. Our saliency model composes of two modules, which are designed for capturing the spatial and temporal saliency stimuli simultaneously. The static saliency network (Sec. III-B) takes single frame as input and outputs static saliency estimates. The dynamic saliency network (Sec. III-C) learns dynamic saliency stimuli from frame pairs and takes static saliency generated by the first module as prior, thus producing the final spatiotemporal saliency maps.

A. Architecture Overview

We start with an overview of our deep video saliency model before going into details below. At a high level, we feed frames of a video into a neural network, and the network successively outputs saliency maps where brighter pixels indicate higher saliency values. The network is trained with video sequences and images and learns generic, spatiotemporal saliency stimuli.

Figure 1 shows the architecture of proposed deep video saliency model. Inspired by classical human visual perception research [11], [12], which suggests that both static and dynamic visual stimuli contribute to dynamic saliency, we design our deep video saliency model two modules, simultaneously considering both the spatial and temporal characteristics of the scene.

The first module is for capturing static saliency stimuli, taking single frame image as input. It adopts fully convolutional networks (FCNs) for generating pixel-wise saliency estimate and utilizes previous excellent pre-trained models on large-scale image datasets. Boosted from rich image saliency benchmarks, this module is efficiently trained for capturing diverse static saliency stimuli of interesting objects. This module is described in detail in Sec. III-B.

The second module takes frame pairs as input and is able to infer dynamic saliency. Additionally, we incorporate static saliency stimuli from the first module as prior knowledge into this network, for more accurate inference and clearer structure. The input frame pair and static saliency stimuli are concatenated and feed into this dynamic saliency network for directly producing final spatiotemporal saliency. This network is trained from both synthetic and real labelled video data (see details in Sec. III-C).

The above architecture has two advantages. First, facilitated by a wide range of image segmentation data and well-trained network in semantic segmentation, generic representations of objects for saliency detection can be easily learned in the first module. Second, static saliency stimuli and saliency




Fig. 2. Illustration of our network for static saliency detection. The network takes single frame image (for example, 224×224) as input, adopting multi-layer convolution networks that transforms the input image to multidimensional feature representation, then applying a stack of deconvolution networks for upsampling the feature extracted from the convolution networks. Finally, a fully convolutional network with 1×1 kernel is used to output of a probability map in the same size as input, in which larger values mean higher saliency values.

priors are naturally encoded into the second module, which makes training deep models easier and yields more accurate prediction.

B. Deep Networks for Static Saliency

A static saliency network takes a single frame image as input and produce a saliency map with the same size of the input. We model this process with a fully convolutional network (FCN). The bottom of this network is a stack of convolutional layers. A convolutional layer is defined on a translation invariance basis and has shared weights across different spatial locations [31]. Each layer input and output in a convolutional network are three-dimensional tensors, called feature maps, with size $h \times w \times c$, where h , w and c are height, width and the feature or channel dimensionality, respectively. The first layer is the image, with pixel size h and w , and three channels. The output feature map is obtained by convolving the input feature map with a linear filter, then adding a bias term. If we denote the input feature map as X , whose convolution filters are determined by the weights W and bias b , then the output feature map is obtained via:

$$f_s(X; W, b) = W *_{\text{s}} X + b, \quad (1)$$

where $*_{\text{s}}$ is the convolution operation with stride s . For improving translation invariance and representation capability, convolutional layers are interleaved with max pooling layers and ReLUs.

Due to the stride of convolutional and pooling layers, the local output feature maps are coarse and downsampled. For upsampling the coarse feature map, multi-layer *deconvolution* (backwards convolution) networks are put on the top of the convolution networks:

$$Y = U_S(F_S(I; \Theta_{\text{con}}); \Theta_{\text{decon}}), \quad (2)$$

where I is the input image; $F_S(\cdot)$ denotes the output feature map generated by the convolutional layers with total stride of S ; $U_S(\cdot)$ denotes the deconvolution layers that upsample the input by a factor of S to ensure the same spatial size of the output Y and the input image I . All the parameters Θ s of convolution and deconvolution layers are learnable.

Finally, on the top of the network, a fully convolutional layer with a 1×1 kernel is adopted for mapping the feature maps Y into a precise saliency prediction map P through a sigmoid activation unit. Due to the utilization of FCN, the network is allowed to operate on input images of arbitrary sizes, and preserves spatial information.

For training, all the parameters Θ s are learned via minimizing a loss function, which is computed as the errors between the probability map and the ground truth. As demonstrated in [47], the use of an asymmetric weighted loss helps greatly in the case of unbalanced data. Considering the numbers of salient and non-salient pixels are usually imbalanced, we compute a weighted cross-entropy loss. Given a training sample (I, G) consisting of an image I with size $h \times w \times 3$, and groundtruth saliency map $G \in \{0, 1\}^{h \times w}$, the network produces saliency probability map $P \in [0, 1]^{h \times w}$. For any given training sample, the training loss on network prediction P is thus given by

$$\mathcal{L}(P, G) = - \sum_{i=1}^{h \times w} ((1 - \alpha)g_i \log p_i + \alpha(1 - g_i) \log(1 - p_i)), \quad (3)$$

where $g_i \in G$ and $p_i \in P$; α refers to ratio of salient pixels in ground truth G .

Figure 2 illustrates the detailed configuration of our deep network for static saliency. The net contains nine layers with weights; the first five are convolutional and the next three are deconvolutional layers. The last layer is a fully convolutional network with a 1×1 kernel and its output is fed to a weighted cross-entropy loss layer which produces a probability map in the same size as the input, in which larger values mean higher saliency values.

We train the proposed architecture in an end-to-end manner. It is commonplace to initialize systems for many of vision tasks with a prefix of a network trained for image classification. This has shown to substantially reduce training time and improve accuracy. During training, our convolutional layers are initialized with the weights in the first convolutional blocks of VGGNet [13], which was originally trained over 1.3 million images of the ImageNet dataset [6]. The parameters of remaining layers are randomly initialized. Then we train




Fig. 3. Illustration of our network for dynamic saliency detection. Successive frame pairs (I_t, I_{t+1}) from real video data or synthesized from existing image datasets (described in Sec. IV), and static saliency prior inferred from our static saliency network, are concatenated and fed into the dynamic network, which has a similar FCN architecture with the static network. The dynamic network captures dynamic saliency stimuli, and considers static saliency prior simultaneously, thus directly generating spatiotemporal saliency estimation.

our network with stochastic gradient descent (SGD) using backpropagation by minimizing the loss in Equ. 3. More details of implementation are described in Sec. V-A.

C. Deep Networks for Dynamic Saliency

Thus far, we have described our deep learning model for detecting static saliency. Next, we describe our spatiotemporal saliency network. As depicted in Figure 3, the network has a similar structure as our static saliency network, which is based on FCN and includes multi-layer convolution and deconvolution nets.

The dynamic network infers dynamic saliency stimuli and explicitly utilizes static saliency prior, thus directly generating spatiotemporal saliency estimates.

The training set consists of a collection of synthetic and real video data, which efficiently utilizes existing large-scale well-annotated image data (described in Sec. IV). More specially, we feed successive pair of frames (I_t, I_{t+1}) and the groundtruth G_t of frame I_t in the training set into this network for capturing dynamic saliency stimuli. Meanwhile, since saliency in dynamic scenes is boosted by both static and dynamic stimuli, the network incorporates the saliency estimate P_t generated by static saliency network as saliency priors indicative of potential salient regions. Thus our dynamic saliency network directly generates spatiotemporal saliency estimates via exploring dynamic saliency stimuli and leveraging static saliency prior from the static saliency network.

Compared with the popular two-stream network structure used in [16], we merge the output of the static network into the dynamic saliency model, which directly produces spatiotemporal saliency results. This architecture brings two advantages. First, the fusion of dynamic and static saliency stimuli is explicitly inserted into the dynamic saliency network, which shows clearer network structure and higher computation efficiency. Second, static saliency, as a strong saliency prior, is introduced for guiding final saliency prediction in dynamic scenes. As demonstrated in [31], incorporating saliency prior

knowledge would make training deep models easier and yield more accurate prediction.

We concatenate frame pair (I_t, I_{t+1}) and static saliency prior P_t in the channel direction, thus generating a tensor \mathbf{I} with size of $h \times w \times 7$. Then we feed \mathbf{I} into our FCN based dynamic saliency network. The first convolution layer is modified as

$$f(\mathbf{I}; W, b) = W_{I_t} * I_t + W_{I_{t+1}} * I_{t+1} + W_{P_t} * P_t + b, \quad (4)$$

where W s represent corresponding convolution kernels; b is bias parameter. During training, stochastic gradient descent (SGD) is employed to minimize the weighted cross-entropy loss described before. After training, given a frame image pair and static saliency prior, the deep dynamic saliency model is able to output final spatiotemporal saliency estimate. More implementation details can be found in Sec. V-A. Qualitative and quantitative study of the effectiveness of our dynamic saliency model is described in Sec. V-C.

IV. SYNTHETIC VIDEO DATA GENERATION

So far, we have described our networks for video saliency detection. We discuss our approach for training our networks for dynamic saliency below. As we demonstrated in Section I, existing video datasets [7], [8], [9], [10] are insufficiently diverse and have very limited scales. As deep learning models are data-driven and have strong learning ability, directly learning deep networks on such video datasets will lead to overfitting. Noticing the gap between the requirement of learning neural networks for video processing and the lack of large-scale, high-quality annotated video data, we propose a technique for synthesizing video data from still frames.

Our video data synthesis strategy, in some way, shares common spirit with existing data augmentation techniques which are widely used in many deep networks and have demonstrated clear contribution for the performance improvement. Classical data augmentation techniques perform simple image transformations, such as mirroring, flipping, rotation




Fig. 4. Illustration of our synthetic video data generation. A synthetic optical flow field (c) is first initialized with considering various motion characters in real video sequences. Via Equ. 6, final optical flow field (f) is generated, which is more smooth and better simulates real motion patterns. According to (f), a synthetic frame image I' and its saliency mask G' are warped from (a) and (b), respectively.

and cropping. However, for the issue of training deep networks with the shortage of large and well-annotated video data, existing data augmentation methods are inapplicable.

Directly deriving video sequences from single image is also impossible. However, our video saliency network takes frame pairs as input, instead of the whole video sequence. That means we can simulate diverse but very short video sequences (only 2 frames in length) via fully utilizing well-labelled large-scale image datasets. Concretely, given a training sample (I, G) from existing image saliency datasets, we wish to generate a pair of frames (I, I') , which present various motion patterns, diverse deformations and smooth transformation, thus being close to real video signal. We start at simulating the correspondence between I' and I , which is easier than directly inferring adjacent frame I' . Let $\mathbf{x} = (x, y)$ denote a point position, the correspondence between I' and I can be represented as an optical flow field $\mathbf{v} = (u, v)$ via:

$$I'(\mathbf{x}) = I(\mathbf{x} + \mathbf{v}(\mathbf{x})). \quad (5)$$

The optical flow field \mathbf{v} directly represents the pixel-level motion information between two neighboring frames. Next we only introduce how to set the vertical displacement u , as the method of generating v is similar.

We model the optical flow on superpixel level as the motion of similar adjacent pixels should present *consistency*. We oversegment I into a group of superpixels \mathcal{R} . According to groundtruth label G , we further divide superpixels \mathcal{R} into foreground superpixels \mathcal{F} and background ones \mathcal{B} , where $\mathcal{R} = \mathcal{F} \cup \mathcal{B}$. For simulating the *diverse motion patterns* of background, we randomly select 10% background regions \mathcal{S} from \mathcal{B} and randomly initialize their motion values u_s (vertical displacement) from $[-d, d]$, where $d = h/10$. The u_s of the other background regions are initialized as zero. The motion patterns of foreground are usually compactness, as the whole foreground regions move more regularly and purposefully compared with background. Beside, the motion between different foreground parts sometimes also present diverse. For example, the whole body of a person go an exact direction but his arms or legs may have different motions. For




Fig. 5. (a) Real images and corresponding saliency groundtruth masks from existing image datasets. (b) Synthetic image examples and saliency masks generated via our method.

this, we first randomly set a value m (from $[-d, d]$) as the main motion patterns of the foreground regions. Then we randomly set vs of foreground regions from $[m - d/10, m + d/10]$ for representing the difference between foreground regions. A visualization of initial optical flow field is presented in Figure 4-a.

A similar process is adopted for generating the initial horizontal motion displacement (v) and we are able to get an initial optical flow \mathbf{v} for I . Next, we propose an energy function for smoothing and propagating the initial optical flow globally, yet preserving the difference between foreground and background in motion patterns. Let the initial motion vector of each superpixel $r_i \in \mathcal{R}$ be denoted as \mathbf{v}_i , the final motion vector $\bar{\mathbf{v}}_i$ is obtained via optimizing the energy function as follows¹:

$$\mathcal{E}(\bar{\mathbf{v}}, \mathbf{v}) = \underbrace{\sum_i \lambda_i (\bar{\mathbf{v}}_i - \mathbf{v}_i)^2}_{\text{Unary Term}} + \underbrace{\sum_{i, i' \in \mathcal{N}} w_{i, i'} (\bar{\mathbf{v}}_i - \bar{\mathbf{v}}_{i'})^2}_{\text{Smooth Term}}. \quad (6)$$

The first term is the unary constraint that each superpixel tends to have its initial motion, while the smooth term gives the interactive constraint that neighboring superpixels have consistent motion patterns when their representative colors are similar. The superpixel neighborhood set \mathcal{N} contains all the spatially adjacent superpixels². The parameter λ is a positive coefficient measuring how much we want to fit the initial motion. Typically, $\lambda = +\infty$ imposes the hard constraint that each region definitely has the initial motion. We define λ :

$$\lambda_i = \begin{cases} 1 & \text{if } r_i \in \mathcal{F} \\ 1 & \text{if } r_i \in \mathcal{S} \\ 10^{-4} & \text{otherwise} \end{cases} \quad (7)$$

¹Here we slightly reuse \mathbf{v} for representing the optical flow vector of superpixel without ambiguity.

²For further improving the motion consistency of background regions, we consider all the selected background regions \mathcal{S} are adjacent in neighboring system \mathcal{N} .

For the seed regions (selected background regions \mathcal{S} and all the foreground regions \mathcal{F}), we expect that they tend to preserve their initial motions; however, for other regions ($\mathcal{B} \setminus \mathcal{S}$), we emphasize more influence on the smooth term thus we can propagate the initial motions from those seed regions.

The weighting function $w_{i,i'}$ in Equ. 6 defines a similarity measure for adjacent superpixels ($r_i, r_{i'} \in \mathbb{N}$):

$$w_{i,i'} = \begin{cases} \exp^{-\|C(r_i) - C(r_{i'})\|^2} & \text{if } r_i, r_{i'} \in \mathcal{F} \\ \exp^{-\|C(r_i) - C(r_{i'})\|^2} & \text{if } r_i, r_{i'} \in \mathcal{B} \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

where $C(r)$ indicates the mean color vector of pixels in superpixel r . We set the weight $w_{i,i'}$ as zero, when two adjacent superpixels are from foreground \mathcal{F} and background \mathcal{B} , respectively. We consider motion consistency inside the foreground and background, while preserve motion difference between foreground and background. Equ. 6 can be efficiently solved by convex optimization and we can obtain a smooth optical flow field \mathbf{v} . As shown in Figure 4, base on \mathbf{v} , we can generate a simulated frame I' and its corresponding annotation G' from (I, G) .

The proposed method is very fast and outputs synthesized video frame pair, optical flow, and pixel-wise annotations simultaneously. The number of samples in existing image segmentation/saliency datasets is ten or hundred order of magnitude larger than in the video segmentation datasets, allowing us to generate enough scenes. For each image sample I of an image dataset, we generate ten simulated frames. Some simulated results can be observed in Figure 5. In our experiments, we use two large image saliency datasets MSRA10K [48] and DUT-OMRON [49], generating more than 150K simulated videos associated with pixel-level annotations and optical flow within 3 hours. Those synthesized video data, combined with real video samples from existing video segmentation datasets, are fed into our CNN based dynamic saliency model, allowing our model to learn general dynamic saliency stimuli and effectively avoid over-fitting.

V. EXPERIMENTAL RESULTS

In this section, we describe our evaluation protocol and implementation details (Sec. V-A), provide exhaustive comparison results over two large datasets (80 videos in total, Sec. V-B), study the quantitative importance of the different components of our system (Sec. V-C), and assess its computational load (Sec. V-D). We compare our method with six top performing saliency detection methods: [29], [50], [15], [37], [14], [38]. In our experiments, all the parameters of our algorithm are fixed to unity.

A. Experimental Setup

1) *Datasets*: We report our performance on two public benchmark datasets: Freiburg-Berkeley Motion Segmentation (FBMS) dataset [7], and Densely Annotated VIdео Segmentation (DAVIS) dataset [10]. The FBMS dataset contains 59 natural video sequences, covering various challenges such as large foreground and background appearance variation,

significant shape deformation, and large camera motion. It is a classical video database that can be used for several vision tasks including video segmentation and saliency detection. The FBMS dataset comes with a split into a training set and a test set, where the training set includes 29 video sequences and the test set has 30 video sequences. We further report our performance on the newly developed DAVIS dataset, which is one of the most challenging video segmentation benchmarks. It consists of 50 video sequences in total, and fully-annotated pixel-level segmentation ground-truth for each frame is available. Those scenes span multiple occurrences of typical challenges (such as occlusions, motion-blur and appearance changes) in video object segmentation.

For testing the proposed video saliency network, we report the performance of our method and other alternatives on the *test set* of FBMS dataset and the *whole* DAVIS dataset.

For training our network, we use two large image saliency datasets: MSRA10K [48] and DUT-OMRON [49]. The MSRA-5000 dataset, comprising of 10K images, is widely used for saliency detection and covers a large variety of image contents - natural scenes, animals, indoor, outdoor, etc. Most of the images have a single salient object. The DUT-OMRON dataset is one of the most challenging image saliency datasets and contains 5172 images with multiple objects with complex structures and high background clutter. All the above datasets contain manually annotated groundtruth saliency/segmentation maps. The video sequences of the whole SegTrackV2 dataset [8] and the training set of the FBMS dataset are also used for training the dynamic saliency network, which include about 3k frame pairs³.

2) *Implementation*: The proposed deep video saliency network has been implemented in MATLAB with the popular Caffe library [51], an open source framework for CNNs training and testing. The source code will be released.

For our static video saliency network, the weights of the first five convolutional blocks are initialized by the VGGNet model [13] trained on ImageNet [6], the other four deconvolutional and the final fully-convolutional networks are initialized from zero mean Gaussian with a standard deviation of 0.01 and the biases are set to 0. Based on this, our network was trained on the MSRA10K [48] and the DUT-OMRON [49] datasets with 100K iterations for saliency detection in static scenes.

Our dynamic video saliency network is also initialized from the VGGNet network. For the first convolutional layer, we use Gaussian initialization due to a different input channel from VGGNet. Benefiting from our video data synthesis approach, we can employ images and annotations from existing saliency segmentation datasets for training our video saliency model. The images and masks from MSRA10K and DUT-OMRON datasets are used to generate more than 1.5×10^5 video slits. Then we combine our simulated video data with real video data (~ 3000 frame pairs) from exiting video segmentation datasets [8], [7] for generating an aggregate video saliency training set. We first pre-train a sub-model of our network without the last two deconvolutional layers

³Due to the number of annotations provided by FBMS is very limited (only 4~6 frames are labeled for each video sequence), we provide extra ~ 500 annotations followed by the original labeling strategy of FBMS.

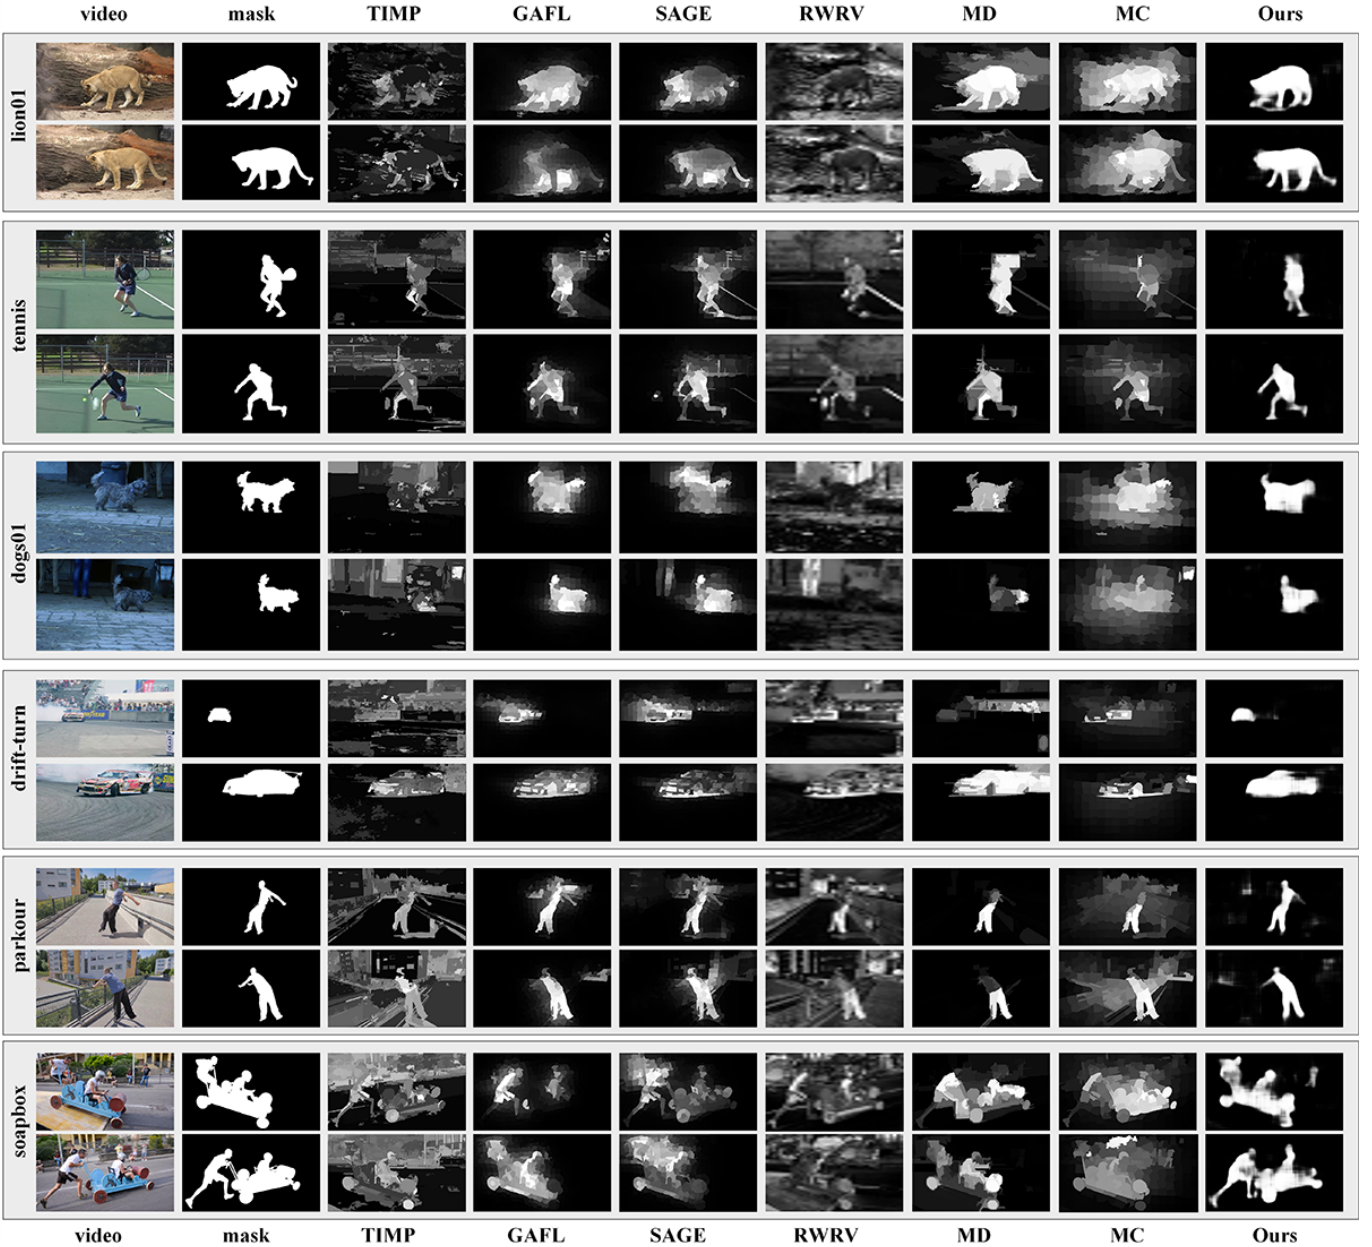


Fig. 6. Qualitative comparison against the state-of-the-art methods on the FBMS dataset videos [7] (*lion01*, *tennis* and *dogs01*), and DAVIS dataset sequences [10] (*drift-turn*, *parkour* and *soapbox*) with pixel-level ground-truth labels. Our saliency method renders the entire objects as salient in complex scenarios, yielding continuous saliency maps that are most similar to the ground-truth.

with 10K iterations, and use this pre-trained model to re-initialize our deep video saliency network. In practice, we find this strategy makes our deep network more easily achieve convergence. Our whole video saliency model is trained for 300K iterations.

For both two networks, we use stochastic gradient descent (SGD) and a polynomial learning policy with initial learning rate of 10^{-7} . The momentum and weight decay are set to 0.9 and 0.0005. The whole training process costs about 40 h on a PC with 3.4 GHz CPU, a TITANX GPU, and 32G RAM.

B. Performance Comparison

To evaluate the quality of the proposed approach, we provide in this section quantitative comparison for perfor-

mance of the proposed method against six top-performing alternatives: saliency via deep feature (MD) [29], saliency via absorbing markov chain (MC) [50], space-time saliency for time-mapping (TIMP) [15], gradient-flow filed based saliency (GAFL) [37], geodesic distance based video saliency (SAGE) [14], and saliency via random walk with restart (RWRV) [38], on test set (30 video sequences) of the FBMS dataset and the whole DAVIS dataset (50 video sequences). The former two methods aim at image saliency while the latter four are designed for video saliency. MD is based on a deep learning model, while others are bottom-up mechanisms.

The results of different methods are achieved by running publicly available implementations with recommended parameter settings provided by the original authors. For fair com-

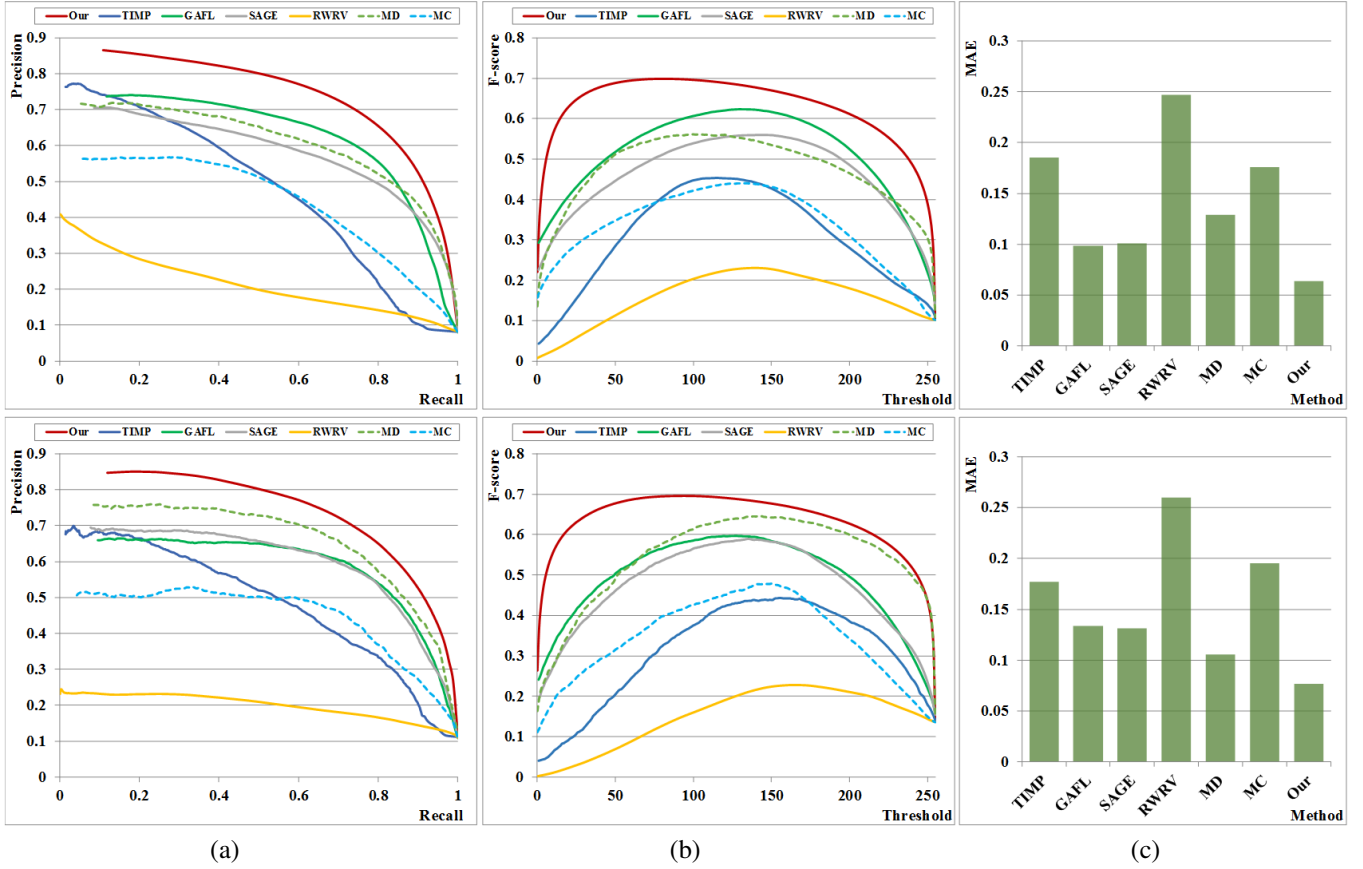


Fig. 7. Comparison with 6 alternative saliency detection methods using the test set of the FBMS dataset [7] (**top**), and the DAVIS dataset [10] (**bottom**) with pixel-level ground-truth: (a) average precision recall curve by segmenting saliency maps using fixed thresholds, (b) F-score, (c) average MAE. Notice that, our algorithm consistently outperforms other methods across different metrics.

parison, all saliency maps generated using different saliency models are normalized into the same range of [0, 255] with the full resolution of original images. From the experimental results, three benefits of our method can be confirmed.

1) *Qualitative Results*: To provide qualitative comparison of the different saliency outputs, we present results of saliency maps generated by our method and six state-of-the-art methods [29], [50], [15], [37], [14], [38] in Figure 6. The top line shows example video frames and the second line shows the ground truth detection results of salient objects.

The image saliency method [50], unsurprisingly, faces difficulties in dynamic scenes, due to the lack of inter-frame information. The video saliency methods [37], [14] generate more visually promising results, but suffer higher computation load (which will be detailed in Sec. V-D) and show relatively weak performance with complex background. As for [29], it is an image saliency model but exhibits competitive performance with above bottom-up video saliency approaches, which demonstrates the power of deep learning model in saliency detection. However, we can observe the proposed algorithm captures foreground salient objects more faithfully in most test cases. In particular, the proposed algorithm yields good performance on some challenging scenarios, even for blurred backgrounds (*lion01*), various object motion patterns (*parkour*) or large shape deformation (*soapbox*). This can be attributed to our video data synthesis, which offers diverse

scene information and rich motion patterns. Based on this, our method is able to learn both static and dynamic saliency stimuli and detects salient moving objects accurately despite similar appearance to the background.

2) *Quantitative Results*: We report quantitative evaluation results on three widely used performance measures: precision-recall (PR) curves, F-measure and MAE. All the evaluation results on a dataset are obtained via averaging the measures over saliency maps of all frames in the dataset.

We first employ precision-recall (PR) curves for performance evaluation. Precision corresponds to the percentage of salient pixels correctly assigned, while recall corresponds to the fraction of detected salient pixels in relation to the ground truth number of salient pixels. For each saliency map, we vary the cutoff threshold from 0 to 255 to generate 256 precision and recall pairs, which are used to plot a PR curve.

The F-measure is the overall performance measurement computed by the weighted harmonic of precision and recall:

$$\text{F-measure} = \frac{(1 + \beta^2) \times \text{precision} \times \text{recall}}{\beta^2 \times \text{precision} + \text{recall}}, \quad (9)$$

where we set $\beta^2 = 0.3$ to weigh precision more than recall as suggested in [52]. For each saliency map, we derive a sequence of F-measure values along the PR-curve with the threshold varying from 0 to 255.

As neither precision nor recall considers the true negative




Fig. 8. Qualitative comparison between our static saliency results and final spatiotemporal saliency results. From top to bottom: input frame images, saliency results via our static saliency network, and spatiotemporal saliency results via our whole video saliency model.

TABLE II
ASSESSMENT OF INDIVIDUAL MODULES AND VARIANTS OF OUR DEEP SALIENCY MODEL ON THE TEST SET OF FBMS DATASET [7] (TOP), AND THE DAVIS DATASET [10] USING MAE. LOWER VALUES ARE BETTER.

aspect	variant	FBMS		DAVIS	
		MAE(%)	Δ MAE(%)	MAE(%)	Δ MAE(%)
	whole model	7.65	-	6.36	-
module	Static model in Sec. III-B	8.19	+0.54	7.17	+0.81
	Dynamic model in Sec. III-C	9.43	+1.78	8.32	+1.96
Training	Only using image data (1.5×10^5)	9.27	+1.62	7.53	+1.17
	Only using video data (0.03×10^5)	24.5	+16.8	23.9	+17.5
	Reduced training data (1×10^5)	9.14	+1.48	7.54	+1.18
	Reduced training data (0.5×10^5)	10.7	+3.08	9.13	+2.77
	Reduced training data (0.1×10^5)	12.8	+5.18	10.9	+4.58
	Reduced training data (0.05×10^5)	13.5	+5.83	12.7	+6.39

saliency assignments, the mean absolute error (MAE) is also introduced as a complementary measure. MAE is defined as the average per-pixel difference between an estimated saliency probability map P and its corresponding ground truth G . Here, P and G are normalized to the interval $[0, 1]$. MAE is computed as

$$\text{MAE} = \frac{\sum_{i=1}^{h \times w} |P(\mathbf{x}_i) - G(\mathbf{x}_i)|}{h \times w}, \quad (10)$$

where h and w refer to the height and width of the input frame image. MAE is meaningful in evaluating the applicability of a saliency model in a task such as object segmentation.

The precision-recall curves of all methods are reported in Figure 7-a. As shown, our method significantly outperforms the state-of-the-art both on the FBMS dataset [7], and the DAVIS dataset [10]. Our saliency method achieves the best precision rates, which demonstrates our saliency maps are more precise and responsive to the actual salient information. The F-scores are depicted in 7-b, in which our model achieves better scores than other methods. Similar conclusions can be drawn from the MAE, as shown in Figure 7-c, where our method achieves the lowest MAE among all compared methods. Overall, our network shows a significant improvement over the current state-of-the-art across challenging datasets.

C. Validation of the Proposed Method

To exhibit more details of our algorithm and objectively evaluate the contribution of different phases in the proposed

saliency model, we report the evaluation of each of the components described in Sec. III and different variants of the proposed saliency model. We experiment on the test set of the FBMS dataset [7], and the DAVIS dataset [10] and measure the performance using MAE.

1) *Ablation study*: We first study the effect of each module of our deep saliency model. In Figure 8, we present qualitative comparison between static saliency from our static network (in Sec. III-B) and final spatiotemporal saliency results from our whole model (in Sec. III-C). It can be observed, due to the lack of dynamic information, the static saliency model faces difficulties distinguishing salient objects from clutter background in dynamic scenes. Via comprehensively utilizing static and dynamic saliency stimuli, our deep video saliency model is able to estimate more accurate spatiotemporal saliency maps.

For quantitatively examining the performance of our static saliency network, we directly use the static saliency maps generated by the static network as final spatiotemporal saliency estimates. We observe decreased performance (7.65 \rightarrow 8.19 on FBMS, 6.36 \rightarrow 7.17 on DAVIS), due to the lack of dynamic saliency stimuli. Similarly, we train a dynamic network without considering static saliency as prior using the same training data. We attribute this to the difficulty of directly capturing dynamic stimuli from two successive frames without any saliency prior or extra motion information. We can draw two important conclusions. First, the fusion of static model and dynamic model improves on both. Second, taking static saliency as prior information makes training the dynamic




Fig. 9. Computational load of our method and the state-of-the-art video saliency methods for processing a 480p video.

model easier and yield more accurate prediction.

2) *Training strategy*: We also explore the effect of different training strategies. We first study the influence of our synthetic video data generation strategy in Sec. IV. We train our deep saliency model only using the synthetics from image data. Although the real video data occupy a small percentage of the training, we can still see a decrease in MAE (7.65→9.27 on FBMS, 6.36→7.53 on DAVIS) when we only use synthetic data. The small performance decrease verifies the effectiveness of our data augmentation technique; on the other hand, it suggests the synthetics should not completely replace the real video data. We further explore the performance of our model only using video data (0.03×10^5 frame pairs). Unfortunately, our model suffers over-fitting due to the high similarities of scenes within same video. This also demonstrates the importance of our synthetic video data generation.

We next study the influence of the amount of training data. When we reduce the amount of training data, we can observe performance decrease. This indicates that the deep-learning model is data-driven. Or, conversely, the increase of training data will lead to improved performance.

D. Runtime Analysis

Here we consider the speed of our saliency method. Our computing platform includes Intel Xeon E7 CPU (12 cores), 64 GB memory and Nvidia Geforce TITAN X GPU (12 GB memory). We do not count I/O time, and do not allow processing multiple images in parallel. The time consumption, of our method compared against other video saliency methods [15], [37], [14], [38] are presented in Figure 9.

From Figure 9 we can learn that, run time efficiency is the major bottleneck for the usability of previous video saliency algorithms, as a substantial amount of time is spent computing motion or edge information. In contrast, our method computes 480p saliency masks in as little as 0.47 seconds, which is much faster than traditional video saliency methods. Our method does not rely on optical flow, edge maps or other pre-computed information, resulting in roughly an order of magnitude faster processing speed.

VI. CONCLUSION

In this work, we have presented a deep learning method for fast video saliency detection using convolutional neural

networks. The proposed deep video saliency model has two modules, namely static saliency network and dynamic saliency network, which are designed for capturing spatial and temporal statistics of dynamic scenes. The saliency estimates from the static saliency network is incorporated in the dynamic saliency network, which enables our method to automatically learn the way of fusing static saliency into dynamic saliency detection and directly produce final spatiotemporal saliency results with less computation load. Furthermore, we proposed a novel data augmentation technique for synthesizing video data from still images, which enables our deep saliency model to learn generic spatial and temporal saliency stimuli and prevents overfitting.

Experimental results on two databases, namely FBMS and DAVIS, have shown that our proposed methods can generate high-quality salience maps. Additionally, our model waives the main computational burdens of previous video saliency models based on optical flow estimation. Our saliency model is very efficient, achieving a processing frame rate of 2fps on a GPU.

REFERENCES

- [1] R. Girshick, “Fast R-CNN,” in *IEEE International Conference on Computer Vision*, 2015, pp. 1440–1448.
- [2] S. Ren, K. He, R. Girshick, and J. Sun, “Faster R-CNN: Towards real-time object detection with region proposal networks,” in *Advances in Neural Information Processing Systems*, 2015, pp. 91–99.
- [3] J. Long, E. Shelhamer, and T. Darrell, “Fully convolutional networks for semantic segmentation,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3431–3440.
- [4] R. Zhao, W. Ouyang, H. Li, and X. Wang, “Saliency detection by multi-context deep learning,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 1265–1274.
- [5] G. Li and Y. Yu, “Deep contrast learning for salient object detection,” *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [6] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, “Imagenet large scale visual recognition challenge,” *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [7] T. Brox and J. Malik, “Object segmentation by long term analysis of point trajectories,” in *European conference on computer vision*, 2010, pp. 282–295.
- [8] F. Li, T. Kim, A. Humayun, D. Tsai, and J. M. Rehg, “Video segmentation by tracking many figure-ground segments,” in *IEEE International Conference on Computer Vision*, 2013, pp. 2192–2199.
- [9] F. Galasso, N. Shankar Nagaraja, T. Jimenez Cardenas, T. Brox, and B. Schiele, “A unified video segmentation benchmark: Annotation, metrics and analysis,” in *IEEE International Conference on Computer Vision*, 2013, pp. 3527–3534.
- [10] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. V. Gool, M. Gross, and A. Sorkine-Hornung, “A benchmark dataset and evaluation methodology for video object segmentation,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [11] A. M. Treisman and G. Gelade, “A feature-integration theory of attention,” *Cognitive psychology*, vol. 12, no. 1, pp. 97–136, 1980.
- [12] P. Mital, T. J. Smith, S. Luke, and J. Henderson, “Do low-level visual features have a causal influence on gaze during dynamic scene viewing?” *Journal of Vision*, vol. 13, no. 9, pp. 144–144, 2013.
- [13] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [14] W. Wang, J. Shen, and F. Porikli, “Saliency-aware geodesic video object segmentation,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3395–3402.
- [15] F. Zhou, S. Bing Kang, and M. F. Cohen, “Time-mapping using space-time saliency,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 3358–3365.
- [16] K. Simonyan and A. Zisserman, “Two-stream convolutional networks for action recognition in videos,” in *Advances in Neural Information Processing Systems*, 2014, pp. 568–576.

- [17] Ç. Bak, A. Erdem, and E. Erdem, "Two-stream convolutional networks for dynamic saliency prediction," *arXiv preprint arXiv:1607.04730*, 2016.
- [18] A. Khoreva, F. Perazzi, R. Benenson, B. Schiele, and A. Sorkine-Hornung, "Learning video object segmentation from static images," *arXiv preprint arXiv:1612.02646*, 2016.
- [19] L. Itti, C. Koch, E. Niebur *et al.*, "A model of saliency-based visual attention for rapid scene analysis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 11, pp. 1254–1259, 1998.
- [20] J. Harel, C. Koch, and P. Perona, "Graph-based visual saliency," in *Advances in Neural Information Processing Systems*, 2006, pp. 545–552.
- [21] T. Judd, K. Ehinger, F. Durand, and A. Torralba, "Learning to predict where humans look," in *IEEE International Conference on Computer Vision*, 2009, pp. 2106–2113.
- [22] F. Perazzi, P. Krähenbühl, Y. Pritch, and A. Hornung, "Saliency filters: Contrast based filtering for salient region detection," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2012.
- [23] M.-M. Cheng, N. J. Mitra, X. Huang, P. H. Torr, and S.-M. Hu, "Global contrast based salient region detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 3, pp. 569–582, 2015.
- [24] W. Wang, J. Shen, L. Shao, and F. Porikli, "Correspondence driven saliency transfer," *IEEE Transactions on Image Processing*, vol. 25, no. 11, pp. 5025–5034, 2016.
- [25] A. Borji and L. Itti, "State-of-the-art in visual attention modeling," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 1, pp. 185–207, 2013.
- [26] A. Borji, M.-M. Cheng, H. Jiang, and J. Li, "Salient object detection: A benchmark," *IEEE Transactions on Image Processing*, vol. 24, no. 12, pp. 5706–5722, 2015.
- [27] Y. Wei, F. Wen, W. Zhu, and J. Sun, "Geodesic saliency using background priors," in *European Conference on Computer Vision*, 2012, pp. 29–42.
- [28] W. Zhu, S. Liang, Y. Wei, and J. Sun, "Saliency optimization from robust background detection," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 2814–2821.
- [29] G. Li and Y. Yu, "Visual saliency based on multiscale deep features," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 5455–5463.
- [30] Y. Tang and X. Wu, "Saliency detection via combining region-level and pixel-level predictions with cnns," in *European Conference on Computer Vision*, 2016, pp. 809–825.
- [31] L. Wang, L. Wang, H. Lu, P. Zhang, and X. Ruan, "Saliency detection with recurrent fully convolutional networks," in *European Conference on Computer Vision*, 2016, pp. 825–841.
- [32] C. Guo, Q. Ma, and L. Zhang, "Spatio-temporal saliency detection using phase spectrum of quaternion fourier transform," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2008, pp. 1–8.
- [33] H. J. Seo and P. Milanfar, "Static and space-time visual saliency detection by self-resemblance," *Journal of vision*, vol. 9, no. 12, pp. 15–15, 2009.
- [34] V. Mahadevan and N. Vasconcelos, "Spatiotemporal saliency in dynamic scenes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 1, pp. 171–177, 2010.
- [35] F. Zhou, S. B. Kang, and M. F. Cohen, "Time-mapping using space-time saliency," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2014.
- [36] Y. Fang, Z. Wang, W. Lin, and Z. Fang, "Video saliency incorporating spatiotemporal cues and uncertainty weighting," *IEEE Transactions on Image Processing*, vol. 23, no. 9, pp. 3910–3921, 2014.
- [37] W. Wang, J. Shen, and L. Shao, "Consistent video saliency using local gradient flow optimization and global refinement," *IEEE Transactions on Image Processing*, vol. 24, no. 11, pp. 4185–4196, 2015.
- [38] H. Kim, Y. Kim, J.-Y. Sim, and C.-S. Kim, "Spatiotemporal saliency detection for video sequences based on random walk with restart," *IEEE Transactions on Image Processing*, vol. 24, no. 8, pp. 2552–2564, 2015.
- [39] J. Charles, T. Pfister, D. Magee, D. Hogg, and A. Zisserman, "Personalizing human video pose estimation," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [40] K. Fragkiadaki, P. Arbelaez, P. Felsen, and J. Malik, "Learning to segment moving objects in videos," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.
- [41] C. Ma, J.-B. Huang, X. Yang, and M.-H. Yang, "Hierarchical convolutional features for visual tracking," in *IEEE International Conference on Computer Vision*, 2015, pp. 3074–3082.
- [42] L. Wang, W. Ouyang, X. Wang, and H. Lu, "Visual tracking with fully convolutional networks," in *IEEE International Conference on Computer Vision*, 2015.
- [43] N. Wang and D.-Y. Yeung, "Learning a deep compact image representation for visual tracking," in *Advances in Neural Information Processing Systems*, 2013, pp. 809–817.
- [44] K. Zhang, Q. Liu, Y. Wu, and M.-H. Yang, "Robust visual tracking via convolutional networks without training," *IEEE Transactions on Image Processing*, vol. 25, no. 4, pp. 1779–1792, 2016.
- [45] H. Nam and B. Han, "Learning multi-domain convolutional neural networks for visual tracking," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [46] Y.-H. Tsai, G. Zhong, and M.-H. Yang, "Semantic co-segmentation in videos," in *European Conference on Computer Vision*, 2016, pp. 760–775.
- [47] M. Mostajabi, P. Yadollahpour, and G. Shakhnarovich, "Feedforward semantic segmentation with zoom-out features," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3376–3385.
- [48] T. Liu, J. Sun, N. N. Zheng, X. Tang, and H. Y. Shum, "Learning to detect a salient object," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2007.
- [49] C. Yang, L. Zhang, H. Lu, X. Ruan, and M.-H. Yang, "Saliency detection via graph-based manifold ranking," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2013.
- [50] B. Jiang, L. Zhang, H. Lu, C. Yang, and M.-H. Yang, "Saliency detection via absorbing markov chain," in *IEEE International Conference on Computer Vision*, 2013.
- [51] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," in *ACM International Conference on Multimedia*, 2014.
- [52] R. Achanta, S. Hemami, F. Estrada, and S. Susstrunk, "Frequency-tuned salient region detection," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.

Wenguan Wang received the B.S. degree in computer science and technology from the Beijing Institute of Technology in 2013. He is currently working toward the Ph.D. degree in the School of Computer Science, Beijing Institute of Technology, Beijing, China. His current research interests include deep learning for salient object segmentation in videos.

Jianbing Shen (M'11-SM'12) is a Professor with the School of Computer Science, Beijing Institute of Technology, Beijing, China. He has published about 60 journal and conference papers such as *IEEE TPAMI*, *IEEE TIP*, *IEEE TVCG*, *IEEE CVPR*, and *IEEE ICCV*. He has also obtained many flagship honors including the Fok Ying Tung Education Foundation from Ministry of Education, the Program for Beijing Excellent Youth Talents from Beijing Municipal Education Commission, and the Program for New Century Excellent Talents from Ministry of Education. His research interests include computer vision and multimedia processing. He serves as an associate editor of *Neurocomputing*.

Ling Shao (M'09-SM'10) is a Professor with the School of Computing Sciences at the University of East Anglia, Norwich, UK. He received the B.Eng. degree in Electronic and Information Engineering from the University of Science and Technology of China (USTC), the M.Sc. degree in Medical Image Analysis and the Ph.D. (D.Phil.) degree in Computer Vision from the University of Oxford. His research interests include Computer Vision, Image/Video Processing, Pattern Recognition and Machine Learning. He has authored/co-authored over 200 papers in refereed journals/conferences such as *IEEE TPAMI*, *TIP*, *TNNLS*, *IJCV*, *ICCV*, *CVPR*, *ECCV*, *IJCAI* and *ACM MM*, and holds over 10 EU/US patents. He is an Associate Editor of *IEEE Transactions on Image Processing*, *IEEE Transactions on Neural Networks and Learning Systems*, *IEEE Transactions on Circuits and Systems for Video Technology*, and other journals. He is a Fellow of the British Computer Society, a Fellow of the IET, and a Life Member of the ACM.